Analysis Facility @ CERN

Enric Tejedor for the SWAN team

https://cern.ch/swan

Analysis Ecosystems Workshop II

May 23rd, 2022

SWAN: the interface of an AF @ CERN

- SWAN: Service for Web-based Analysis
- CERN's Jupyter notebook service
  - Created in 2016
  - Used by 200-250 people daily
- Jupyter interface + federation of CERN services → added value!
  - Software (CVMFS)
  - Storage (EOS, CERNBox)
  - Computing resources (GPU, Spark, HTCondor)
- Platform for physics analysis: supports both *single-node* and *distributed* analysis

Storage: EOS, CERNBox

- Find the data you need for your analysis
  - EOS: experiment repositories (/eos/atlas, /eos/cms, …), projects, open data
  - CERNBox as home directory, sync & share

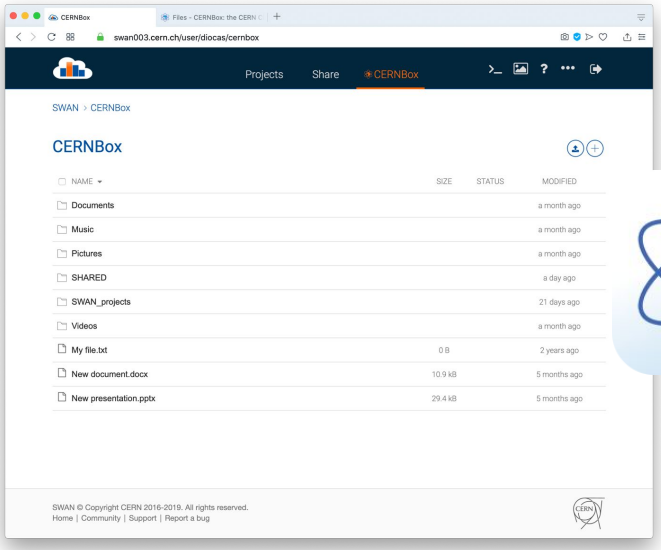

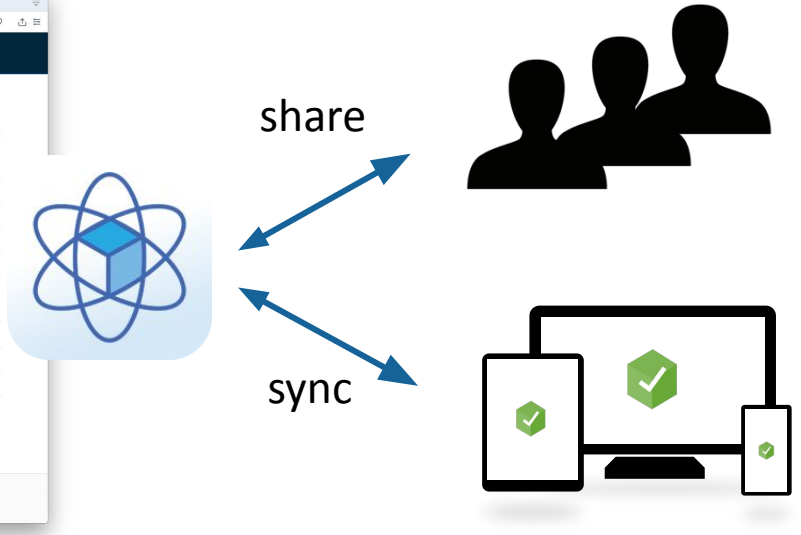

Software: CVMFS

- Find the software you need for your analysis
  - CVMFS: LCG releases (and soon CMSSW, FCC)
  - EOS: custom software environment

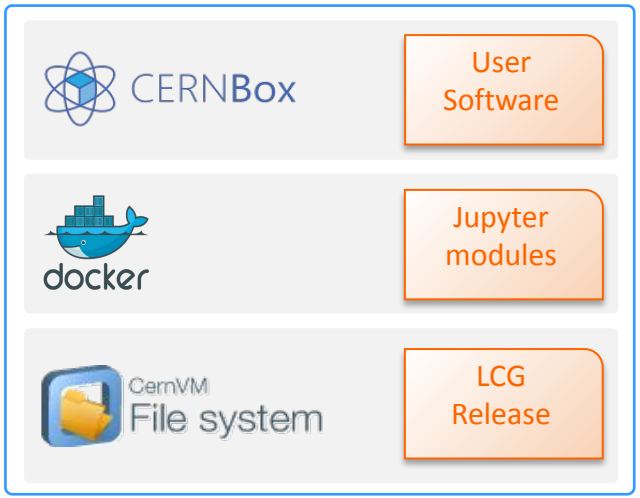

custom user env (optional)

thin layer (not user defined)

main software source

GPUs

- SWAN allows to attach a GPU to a user session
  - Feature of the new SWAN k8s deployment (https://swan-k8s.cern.ch)
  - ~10 GPUS (Tesla T4 and V100)
- The GPUs are used interactively
  - When starting their session, the user selects a CUDA software stack and gets a GPU
  - GPU-enabled packages (e.g. tensorflow, PyTorch) can then be used in a notebook and offload to the GPU by default

```
In [1]: import tensorflow as tf

        tf.debugging.set_log_device_placement(True)

        # Create some tensors
        a = tf.constant([[1.0, 2.0, 3.0], [4.0, 5.0, 6.0]])
        b = tf.constant([[1.0, 2.0], [3.0, 4.0], [5.0, 6.0]])
        c = tf.matmul(a, b)

        Executing op MatMul in device /job:localhost/replica:0/task:0/device:GPU:0
```

Spark

- SWAN is connected to the Spark clusters at CERN
  - Physical: ~3800 cores, some dedicated
  - Virtual: ~250 cores, on demand (kubernetes)
- Jupyter extensions available to:
  - Connect to a certain cluster
  - Monitor the execution

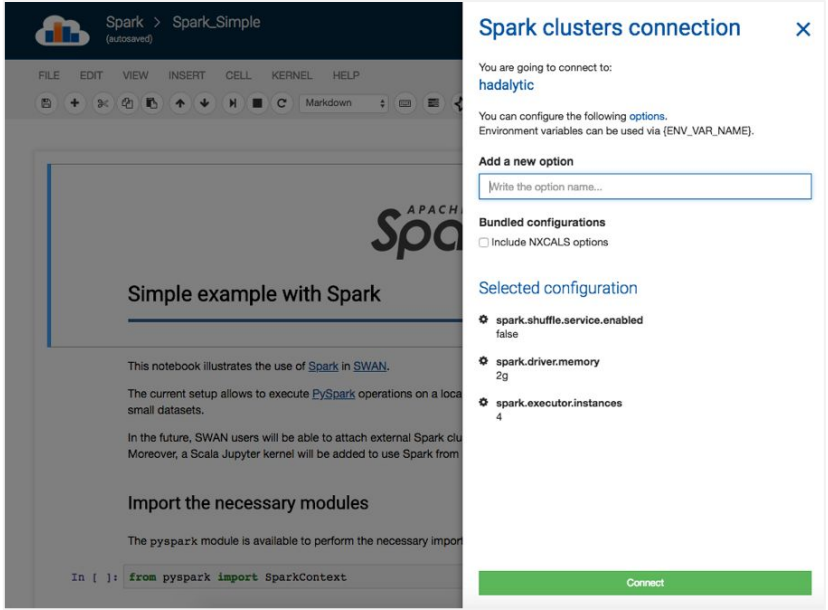

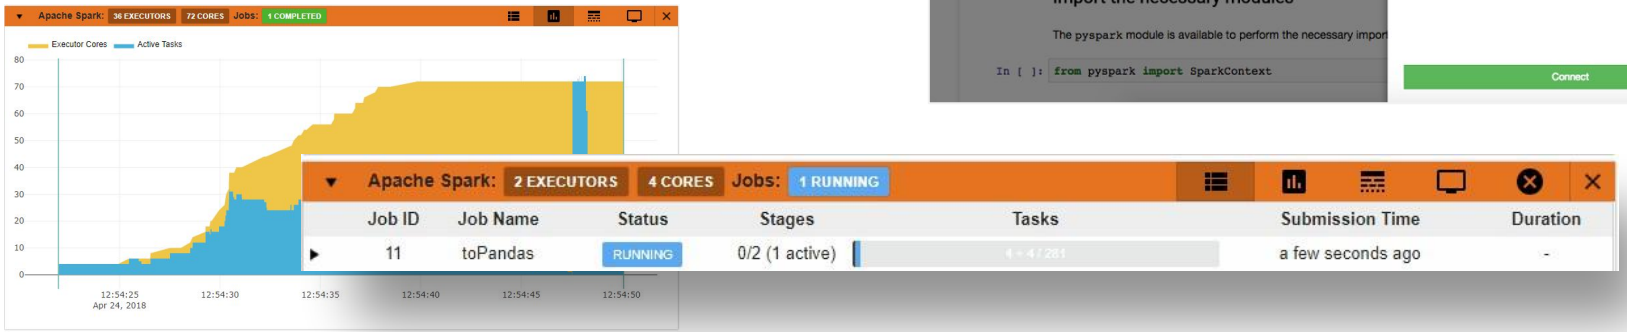

HTCondor

- Goal: leverage HTCondor pools at CERN from SWAN
  - Up to ~175k cores in shared pools at CERN – limited by the quotas assigned depending on experiment affiliation
  - Already used for analysis
- Batch submission: already supported
  - Condor packages available on CVMFS
- Interactive usage: in pilot phase
  - Collaboration with Batch Service@CERN
  - Dask packages available on CVMFS
  - Will be exposed to users when migration to JupyterLab is finished (Q2-Q3 2022)

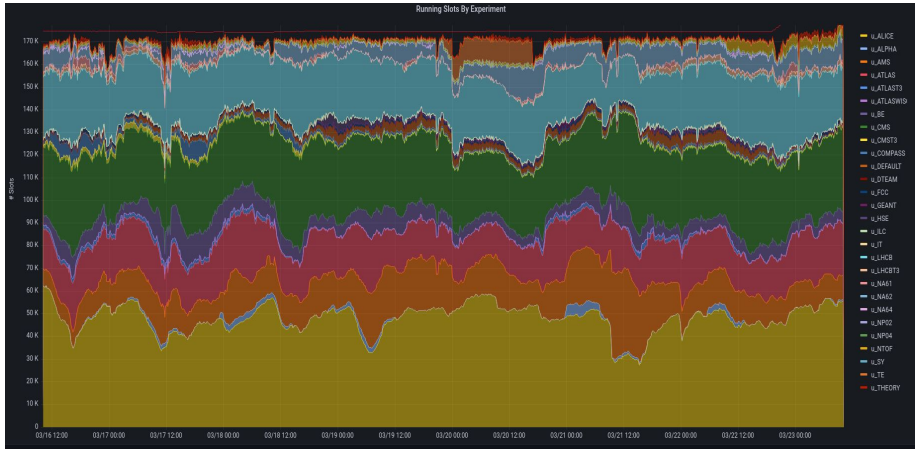

SWAN + HTCondor for interactive analysis

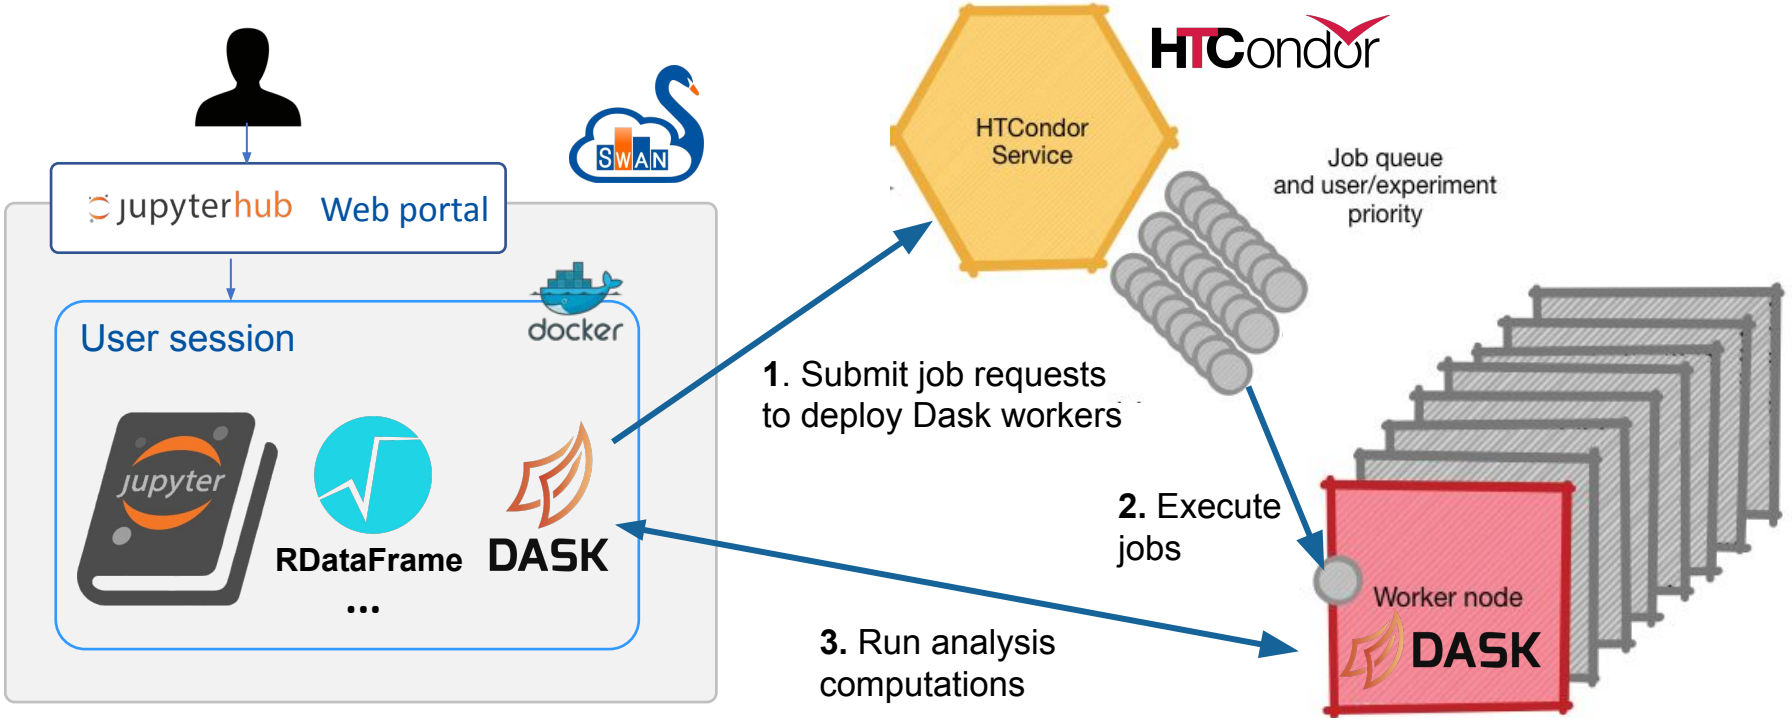

REANA Reusable Analysis Platform

- Born within analysis preservation and reuse context
- Run computational workflows on containerised clouds
  - multiple workflow systems (CWL, Serial, Snakemake, Yadage)
  - multiple container technologies (Docker, Singularity)
  - multiple compute backends (Kubernetes, HTCondor, Slurm)

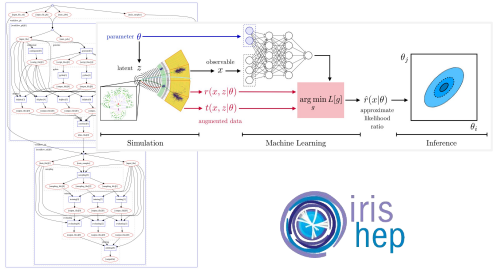

Example: MadMiner ML workflows

reana @BNL @NYU @ND

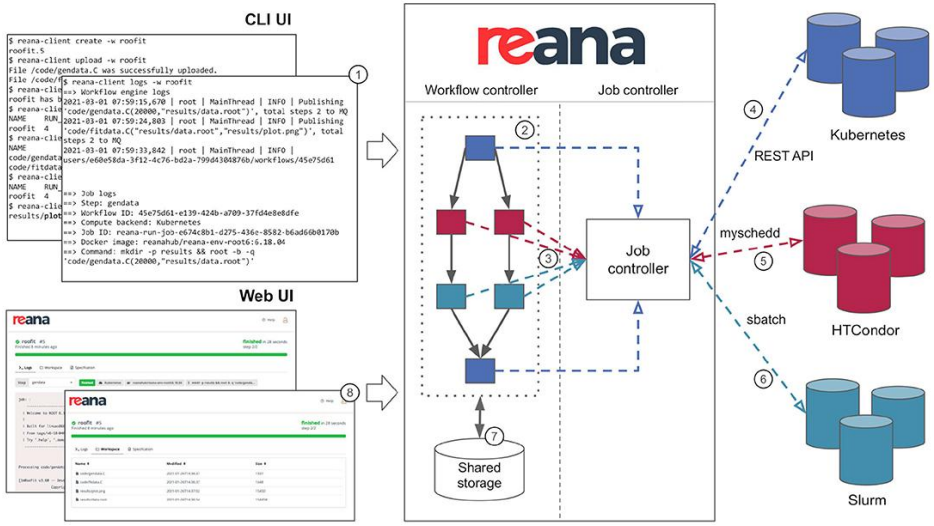

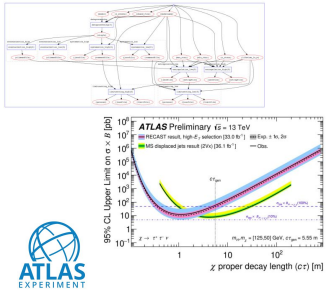

Example: ATLAS search reinterpretations

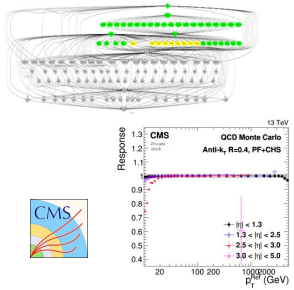

Example: CMS jet energy resolution and corrections

[More info here]

Backup slides

Context and motivation

- HL-LHC needs are pushing us to build **modern Analysis Facilities**
  - Traditional batch processing
  - Interactive computing on big datasets, with new interfaces (Jupyter)
- An AF should facilitate access to:
  - Software
  - Storage (+ sharing)
  - Computing resources (elastic)
- Ongoing effort to provide an AF @ CERN
  - Interdepartmental collaboration (EP, IT)
  - In contact with Analysis Facility WG
- Build on what already exists whenever possible!

SWAN unique users in 2021 (by experiment)

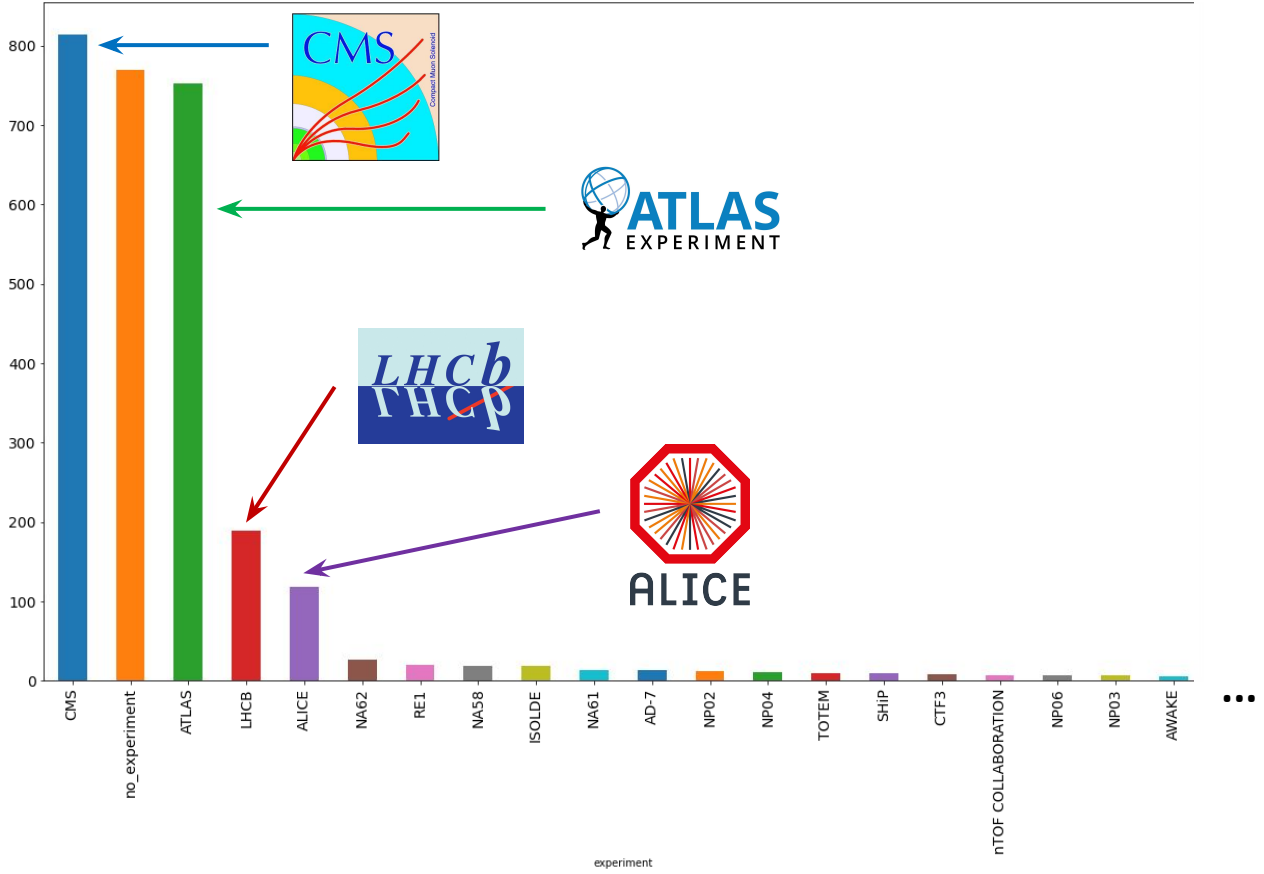

SWAN's building blocks

ScienceBox: installable SWAN

- SWAN can be installed on premises thanks to ScienceBox
  - Packaged SWAN, CVMFS, EOS (and soon CERNBox)
  - https://sciencebox.web.cern.ch
- Two alternatives for installation
  - Single-node: for testing, minikube
  - Multi-node: for production, kubernetes Helm charts
- Successfully deployed outside CERN
  - Aarnet, JRC, education and outreach projects
  - In progress/discussion: WUR, Purdue university (CMS Tier 2)
- In sync with CERN's production SWAN
  - Will benefit too from the integration with Dask and resource managers (HTCondor, kubernetes)

Spark clusters

Cluster Name	Configuration	Primary Usage
analytix	46 nodes (Cores – 1956, Mem – 24.4 TB, Storage – 17.5 PB)	General Purpose
nxcals	38 nodes (Cores – 1820, Mem – 17 TB, Storage – 13 PB)	Accelerator logging (NXCALS) project dedicated cluster
Cloud containers	OpenStack project, Spark-as-a-Service, CPU-optimized (Cores 256, Mem – 2 TB, Storage – EOS) + possibly more	General Purpose Compute ONLY

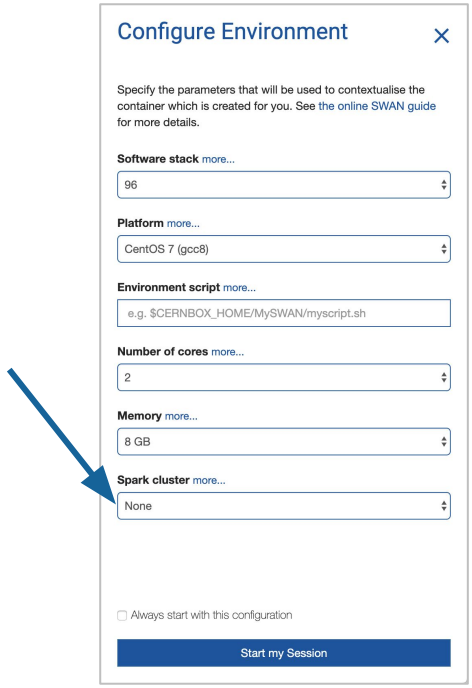

Spark Connector

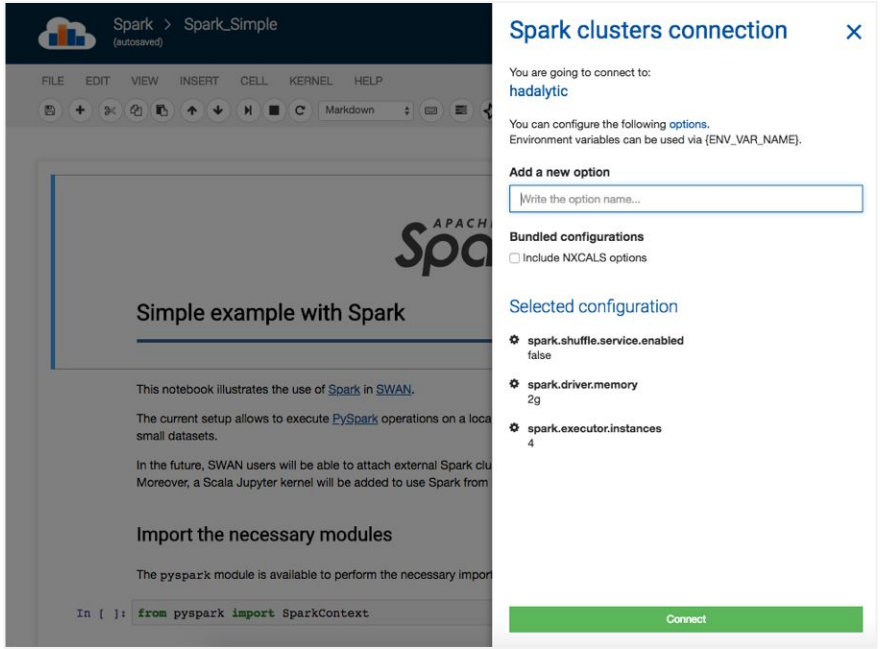

- **Spark Connector** – handling the spark configuration complexity
  - User is presented with Spark Session (Spark) and Spark Context (sc)
  - Ability to bundle configurations specific to user communities
  - Ability to specify additional configuration

Spark Monitor

- **Spark Monitor** – Jupyter notebook extension
  - For live monitoring of spark jobs spawned from the notebook
  - A graph showing number of active tasks & executor cores vs time
  - A timeline which shows jobs, stages, and tasks

Google Summer of Code

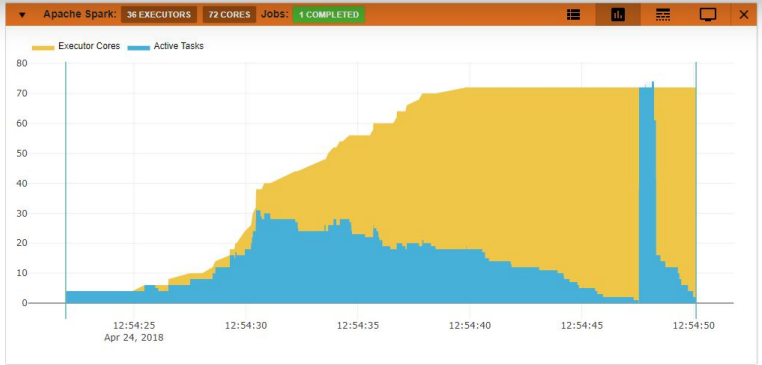

HTCondor @ CERN